

Strong Chain Rules for Min-Entropy under Few Bits Spoiled

Maciej Skórski

IST Austria

Email: mskorski@ist.ac.at

Abstract—It is well established that the notion of min-entropy fails to satisfy the *chain rule* of the form $H(X, Y) = H(X|Y) + H(Y)$, known for Shannon Entropy. The lack of a chain rule causes a lot of technical difficulties, particularly in cryptography where the chain rule would be a natural way to analyze how min-entropy is split among smaller blocks. Such problems arise for example when constructing extractors and dispersers.

We show that any sequence of variables exhibits a very strong strong block-source structure (conditional distributions of blocks are nearly flat) when we *spoil few correlated bits*. This implies, conditioned on the spoiled bits, that *splitting-recombination properties* hold. In particular, we have many nice properties that min-entropy doesn't obey in general, for example strong chain rules, “information can't hurt” inequalities, equivalences of average and worst-case conditional entropy definitions and others.

Quantitatively, for any sequence X_1, \dots, X_t of random variables over an alphabet \mathcal{X} we prove that, when conditioned on $m = t \cdot O(\log \log |\mathcal{X}| + \log \log(1/\epsilon) + \log t)$ bits of auxiliary information, all conditional distributions of the form $X_i|X_{<i}$ are ϵ -close to be nearly flat (only a constant factor away). The argument is combinatorial (based on simplex coverings).

This result may be used as a generic tool for *exhibiting block-source structures*. We demonstrate this by reproving the fundamental converter due to Nisan and Zuckermann (*J. Computer and System Sciences*, 1996), which shows that sampling blocks from a min-entropy source roughly preserves the entropy rate. Our bound implies, only by straightforward chain rules, an additive loss of $o(1)$ (for sufficiently many samples), which qualitatively meets the first tighter analysis of this problem due to Vadhan (CRYPTO'03), obtained by large deviation techniques.

Keywords—chain rule, min-entropy, spoiling knowledge, block sources, local extractors

I. INTRODUCTION

A. Strong vs Weak Entropy Chain Rules

One of the most useful properties of Shannon entropy is the chain rule, showing how entropy splits between distributions

$$H_1(X|Y) = H(X, Y) - H(Y). \quad (1)$$

The notion of min-entropy, very important for cryptography [Sh11], fails to satisfy this property [IS13]. In the lack of a chain rule, a much weaker one-sided bound (e.g. of the form $\tilde{H}_\infty(X|Y) \geq H_\infty(X) - H_0(Y)$, where \tilde{H}_∞ is an appropriate extension of min-entropy to conditional distributions) is sometimes used [DORS08], which we address as a *weak chain rule*¹.

¹In leakage-resilient cryptography such bounds are simply called chain rules. In this paper we discuss chain rules in a strong sense.

B. Need for Strong Chain Rules

While the weak chain rule suffices for many applications related to bounded leakage [DORS08; DP08] (where Y is leakage much shorter than the amount of min-entropy in X), it is insufficient where one needs to estimate how entropy is distributed among blocks. In these settings, one would like to argue that (roughly) either X or $Y|X$ has high min-entropy if the joint min-entropy (of (X, Y)) is high. Examples of such problems are randomness extraction in the bounded storage model [Cac97], constructions of dispersers [BRWS06], or oblivious transfer protocols [DFRSS07].

C. Our Contribution and Related Works

Although the chain rule fails in general, we show that it is true conditioned on *few spoiled bits*. We actually show more, that (locally, conditioned on auxiliary bits) a very strong *block-source structure* exists. Namely, each block is nearly flat given previous blocks. Informally, the theorem reads as follows

Theorem (Informal: exhibiting flat block-source structures). *For any sequence $X = X_1, \dots, X_t$ of correlated random variables each over \mathcal{X} and any ϵ there exists auxiliary information S of length m bits such that*

- (a) *S is short: $m = t \cdot O(\log \log |\mathcal{X}| + \log \log(1/\epsilon) + \log t)$*
- (b) *Conditioned on S , conditional block distributions $\Pr[X_i|X_i = x_i, X_{i-1} = x_{i-1}, \dots, X_1 = x_1]$ are nearly flat (ϵ -close to a probability distribution whose values differ by a constant factor).*

The formal statement is given in Theorem 1. For cryptographic applications, $\log \log(\epsilon^{-1})$ is pretty much a small constant (typically $\epsilon = 2^{-100}$). Also, for sources with super-logarithmic entropy per block, that is when $H_\infty(X) \gg t \log \log |\mathcal{X}|$, and the number of blocks t growing not too fast, e.g. $t = \log^{O(1)}(|\mathcal{X}|)$, the error term is of a smaller order than the entropy. Under these mild assumptions, conditioned on the partition generated by auxiliary bits, we conclude many nice properties that fail in general. Examples are chain rules, “conditioning only decreases entropy” properties, equivalences of conditional entropy defined in different ways and others.

1) Our Tools:

a) *Spoiling Knowledge*: The *spoiling knowledge* technique is essentially about finding auxiliary information that increases entropy, and was introduced in [BBCM95]. We use the same idea to force block distributions to be nearly flat.

b) *Covering techniques (combinatorial geometry)*: In order to construct a good “spoiling”, we consider the logarithm of the chaining identity $p_{X_1 \dots X_t}(\cdot) = \prod_{j=1}^t p_{X_j | X_{j-1} \dots X_1}(\cdot)$ which represents the “surprise” of the total distribution as a sum of “next-block surprises” of the form $r_i(\cdot) = -\log p_{X_i | X_{i-1} \dots X_1}(\cdot)$. It follows that the vector of all r_i (for $i = 1, \dots, t$) lie in a $(t-1)$ -dimensional simplex of edge roughly $O(t \log |\mathcal{X}|)$. Our partition is obtained from *coverings*, as for all x in the same part $r_i(\cdot)$ is roughly constant (when the radius is small enough). This approach can be likely optimized (we use a crude bound on the covering number). Also the lower bounds on the necessary number of spoiled buts are possible, by considering packings instead of coverings. We defer this discussion to the full version.

D. Related works

As far as we know, the presented result on spoiling min-entropy is new. The Nisan-Zuckerman lemma, discussed in this paper as an application, was analyzed by Vadhan [Vad03] and recently by Bellare [BKR16]. These results study min-entropy present in random chunks of a larger source, and don’t offer tools for *splitting entropy* in any source *deterministically* (particularly for a small number of blocks), as we do.

E. Applications

The important result due to Nisan and Zuckerman [NZ96], improved later by Vadhan [Vad03] states that sampling from a given source of high min-entropy rate α yields a source of a comparable entropy rate β . This fact is a crucial step in constructions of so called *local extractors*, that extract randomness parsing only a part of input. For a while, only a lossy bound $\beta \approx \alpha / \log(1/\alpha)$ was known. The reason was precisely the lack of a chain rule for min-entropy. As observed by Cachin [Cac97] the proof for Shannon entropy (a less interesting case) is straightforward and follows basically by a splitting-recombining argument, which uses a chain rule in both directions. We demonstrate by our technique that (surprisingly) a very effective splitting-recombining approach actually works, and achieves $\beta = \alpha - o(1)$ in a very straightforward way. This matches the bound due to Vadhan. Concretely, if the original source is a sequence of t blocks over an alphabet \mathcal{X} and we take ℓ samples, then $\beta = \alpha - \text{err}_{\text{Spoil}} - \text{err}_{\text{Samp}}$ where the losses due to chain rules and sampling equal, respectively

$$\begin{aligned} \text{err}_{\text{Spoil}} &= O(\log \log |\mathcal{X}| + \log \log(1/\epsilon) + \log t) / \log |\mathcal{X}| \\ \text{err}_{\text{Samp}} &= O(\sqrt{\ell^{-1} \log(1/\epsilon)}). \end{aligned}$$

In particular β converges to α when the block length is $\log |\mathcal{X}| = \log \omega(1)(1/\epsilon)$ and $\log |\mathcal{X}| = \omega(1) \cdot t$. For more details see Theorem 2 in Section V. For the discussed result our bounds converge slightly slower than Vadhan’s bounds derived by large deviation techniques. However, our spoiling technique can be used also for small number of samples.

F. Organization

In Section II we explain necessary notions and notations. Auxiliary facts that will be needed are discussed in Section III.

In Section IV we prove the main result. Applications to the bounded storage model are discussed in Section V. We conclude the work in Section VI.

II. PRELIMINARIES

A. Basic Notation

For any random variables X_1, X_2 by $p_{X_1 | X_2}$ we denote the distribution of X_1 conditioned on X_2 , that is $p_{X_1 | X_2}(x_1, x_2) = \Pr[X_1 = x_1 | X_2 = x_2]$. Throughout this paper, all logarithms are taken to base 2. For any sequence of random variables $X = X_1, \dots, X_n$ we denote $X_{<i} = X_1, \dots, X_{i-1}$, $X_{\leq i} = X_1, \dots, X_i$ and more generally for any subset $I \subset \{1, \dots, n\}$ we put $X_I = X_{i_1} X_{i_2} \dots X_{i_{m-1}} X_{i_m}$ where $i_1 < i_2 < \dots < i_m$ are all elements of I .

B. Distances, Entropies

In the definitions below \mathcal{X} is an arbitrary finite set.

Definition 1 (Statistical Distance). *For two random variables X, Y on \mathcal{X} by the statistical distance (total variation) we mean*

$$d_{TV}(X; Y) = \frac{1}{2} \sum_{x \in \mathcal{X}} |\Pr[X = x] - \Pr[Y = x]|$$

Definition 2 (Shannon Entropy). *The Shannon entropy of a random variable X on \mathcal{X} equals*

$$H_1(X) = - \sum_x \Pr[X = x] \log \Pr[X = x].$$

Definition 3 (Min-Entropy). *The min-entropy of a random variable X on \mathcal{X} equals*

$$H_\infty(X) = - \log \max_{x \in \mathcal{X}} \Pr[X = x].$$

Definition 4 (Conditional Min-Entropy [DORS08]). *Let X, Y be random variables over \mathcal{X} and \mathcal{Y} respectively. The worst-case min-entropy of X conditioned on Y equals*

$$H_\infty(X|Y) = \min_{y \in \mathcal{Y}} H_\infty(X|Y=y).$$

The average min-entropy of X conditioned on Y equals

$$\tilde{H}_\infty(X|Y) = - \log \left(\mathbb{E}_{y \leftarrow Y} 2^{-H_\infty(X|Y=y)} \right).$$

Remark 1. *The averaged notion is slightly weaker, but has better properties and actually better suits applications [DORS08].*

The notion of smooth entropy is more accurate than min-entropy because quantifies entropy up to small perturbations in the probability mass.

Definition 5 (Smooth Min-Entropy [RW04]). *The ϵ -smooth min-entropy of a random variable X on \mathcal{X} is defined as*

$$H_\infty^\epsilon(X) = \max_{X': d_{TV}(X, X') \leq \epsilon} H_\infty(X')$$

where the maximum is over all random variables X' on \mathcal{X} . In other words, X has at least k bits of smooth min-entropy if there is X' of min-entropy at least k and ϵ -close to X .

C. Randomness Extractors

Below we recall the definition of seeded extractors

Definition 6 (Randomness Extractor [NZ96]). *We say that a function $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a (k, ϵ) -extractor if and only if*

$$d_{TV}(\text{Ext}(X, U_d), U_d; U_m, U_d) \leq \epsilon$$

for any X on $\{0, 1\}^n$ with min-entropy at least k .

D. Samplers

Averaging samplers are procedures which sample points, within a given domain, that are distributed enough random to approximate every function. It turns out that there exist good averaging samples, using much less auxiliary randomness than necessary to generate independent points. For our applications we consider samplers that approximate averages from below.

Definition 7 (Averaging Samplers [Vad03]). *A function $\text{Samp} : \{0, 1\}^d \rightarrow [t]^\ell$ is a (μ, θ, γ) -averaging sampler if for every function $f : [t] \rightarrow [0, 1]$ with average value $\frac{1}{\ell} \sum_{i=1}^\ell f(i) \geq \mu$ it holds that*

$$\Pr_{i_1, \dots, i_\ell \leftarrow \text{Samp}(U_r)} \left[\frac{1}{\ell} \sum_{i=1}^\ell f(i) \geq \mu - \theta \right] \leq 1 - \gamma.$$

Lemma 1 (Optimal Averaging Samplers (Nonconstructive) [Vad03]). *For the setting in Definition 7 there is an averaging sampler which for any ℓ' such that $\Omega(\mu\theta^{-2} \log(1/\gamma)) \leq \ell' \leq \ell$ produces ℓ' distinct samples and uses $d = \log(t/\ell') + \log(1/\gamma) + 2 \log(\mu/\theta) + \log \log(1/\mu) + O(1)$ random bits.*

III. AUXILIARY LEMMAS

The following lemma is essentially the *information can't hurt* principle, well known for Shannon entropy, stated for the notion of min-entropy. It says that extra information can only reduce the uncertainty, and comes up often in cryptography.

Lemma 2 (Conditioning only reduces entropy [IS13]). *For any X, Y, Z we have $\tilde{H}_\infty(X|Y, Z) \leq \tilde{H}_\infty(X|Y)$.*

It is well known (see, e.g., [BST03]) that when each block has certain min-entropy entropy conditioned on previous blocks, the total entropy grows linearly with the number of blocks (as expected). The lemma is stated without a proof (it follows directly from definitions). We stress, however, that using the *worst-case* notion of conditional entropy is crucial.

Lemma 3 (Min-entropy from block sources). *For any X, Y, Z we have $H_\infty(X, Y|Z) \geq H_\infty(X|Y, Z) + H_\infty(Y|Z)$.*

Another lemma, well known in folklore, quantifies the intuition that conditioning on “large” events (not too surprising information) doesn't decrease entropy too much.

Lemma 4 (Leakage lemma for min-entropy). *For any random variable X and any events A, B we have $H_\infty(X|B, A) \geq H_\infty(X|B) - \log(1/\Pr(A|B))$.*

Proof. The first part of the lemma follows because we have $\Pr[X = x|A, B] = \Pr[X = x, A \cap B] / \Pr[A \cap B] \leq \Pr[X = x, B] / (\Pr[B] \cdot \Pr[A|B])$ for every x . \square

IV. MAIN RESULT

A. Chain Rule

Theorem 1 (Exhibiting block structures with few bits spoiled). *Let \mathcal{X} be a fixed alphabet and $X = (X_1, \dots, X_t)$ be a sequence of (possibly correlated) random variables each over \mathcal{X} . Then for any $1 > \epsilon > 0$ and $\delta > 0$ there exists a collection \mathcal{B} of disjoint sets on \mathcal{X}^t such that*

(a) \mathcal{B} can be indexed by a small number of bits, namely

$$\log |\mathcal{B}| = t \cdot O(\log \log |\mathcal{X}| + \log \log(\epsilon^{-1}) + \log(t/\delta))$$

(b) \mathcal{B} almost covers the domain

$$\sum_{B \in \mathcal{B}} p_X(B) \geq 1 - \epsilon$$

(c) Conditioned on members of \mathcal{B} , block distributions $X_i|_{X_{<i}}$ are nearly flat.

$$\forall x, x' \in B : 2^{-O(\delta)} \leq \frac{p_X(x_i|x_{<i})}{p_X(x'_i|x'_{<i})} \leq 2^{O(\delta)}.$$

for every $B \in \mathcal{B}$ and $i = 1, \dots, t$.

In some applications (see for example Section V) it is convenient to work with parts that are not too small. By substituting $\epsilon := \frac{\epsilon}{2}$, and deleting from \mathcal{B} all members of smaller than $\frac{\epsilon}{2|\mathcal{B}|}$ (their mass is at most $\frac{\epsilon}{2}$), we obtain

Remark 2 (Getting rid of tiny parts). *In Theorem 1, we may assume $p_X(B) = \Omega(\epsilon \cdot |\mathcal{B}|^{-1})$ for every $B \in \mathcal{B}$.*

Corollary 1 (Conditional entropies under few bits spoiled). *Under the assumptions of Theorem 1, for every $B \in \mathcal{B}$ for every index i and for every set $I \subset \{1, \dots, i-1\}$ we have*

(a) *The chain rule for min-entropy*

$$H_\infty(X_i|X_I, B) = H_\infty(X_i, X_I|B) - H_\infty(X_I|B) \pm O(\delta).$$

(b) *The average and worst-case min-entropy almost match*

$$\tilde{H}_\infty(X_i|X_I, B) = H_\infty(X_i|X_I, B) \pm O(\delta).$$

Proof of Corollary 1. Fix any subset $I \subset \{1, \dots, i-1\}$ of size m . By Theorem 1 for every x and B we have

$$\Pr[X_{\leq i} = x_{\leq i}|B] = 2^{\pm O(\delta)} \Pr[X_{<i} = x_{<i}|B]. \quad (2)$$

Let $J = \{1, \dots, i-1\} \setminus I$. Taking the sum of Equation (2) over x such that $x_I = x'_I$ and $x_i = x''$ are fixed but x_J varies we obtain

$$\Pr[X_i = x'', X_I = x'_I|B] = 2^{\pm O(\delta)} \Pr[X_I = x'_I|B].$$

This implies

$$H_\infty(X_i|X_I, B) = H_\infty(X_i, X_I, B) - H_\infty(X_I, B) \pm O(\delta)$$

$$\tilde{H}_\infty(X_i|X_I, B) = H_\infty(X_i|X_I, B) \pm O(\delta)$$

which finishes the proof. \square

Proof of Theorem 1. Let p be the joint distribution of X_1, \dots, X_n . For any i denote

$$p_i(x_i | x_{i-1}, \dots, x_1) = p_{X_i | X_{<i}}(x_i, \dots, x_1). \quad (3)$$

and let the “surprise” of the bit x_i be

$$r^i(x) = -\log p_i(x_i, \dots, x_1).$$

Note that $p(x_1, \dots, x_n) = \prod_{i=1}^n p_i(x_i, x_{i-1}, \dots, x_1)$ and therefore, denoting $x = (x_1, \dots, x_n)$, we obtain

Claim 1 (Decomposing surprises). *We have*

$$\sum_{i=1}^t r^i(x) = -\log p(x). \quad (4)$$

The next claim follows by a simple Markov-type argument.

Claim 2 (Significant probabilities). *There exists a set $A \subset \mathcal{X}^t$ of probability $1 - \epsilon$ such that $p(x) \geq \frac{\epsilon}{|\mathcal{X}|^t}$ for all $x \in A$.*

Denoting $x = x_1, \dots, x_n$, we have

$$\forall x \in A: -\log p(x_1, \dots, x_n) \leq t \log |\mathcal{X}| + \log(1/\epsilon). \quad (5)$$

The claim below follows from Equation (3) and Equation (4)

Claim 3 (Surprises live in a simplex). *We have $r^i(x) \geq 0$ for $i = 1, \dots, t$ and $\sum_{i=1}^t r^i(x) \leq t \log |\mathcal{X}| + \log(1/\epsilon)$, for all points $x \in A$.*

Claim 4 (Simplex coverings imply a chain rule). *If the simplex with side length $t \log |\mathcal{X}| + \log(1/\epsilon)$ can be covered by N balls of radius R in the ℓ_∞ norm, then the theorem holds with $|\mathcal{B}| = N$ and $\delta = R$.*

Proof of ?? 4. Let $C \subset \mathbb{R}^t$, $|C| = N$, be the set of the centers of the covering balls. Let S be the function which assigns to every point $x \in A$ (where A is defined in ?? 2) the point $z \in C$ closest to the vector $(r_1(x), \dots, r_t(x))$ in the ℓ_∞ -norm. Fix any z and let $B_z = \{x : S(x) = z\}$. By the properties of the covering, for $i = 1, \dots, t$ we obtain

$$\forall x \in B: |r^i(x) - z_i| \leq R \quad (6)$$

In particular, the surprises for any two points in B are close

$$\forall x, x' \in B: |r^i(x) - r^i(x')| \leq 2R, \quad i = 1, \dots, t. \quad (7)$$

Let p_B be the conditional probability of p given B . Denote by r_B^i the surprise of the i -th bit given previous bits and conditioned on B , that is

$$\begin{aligned} r_B^i(x) &= -\log \Pr[X_i = x_i | X_{i-1} = x_{i-1}, \dots, X_1 = x_1, B] \\ &= -\log p_B(x_i | x_{i-1}, \dots, x_1) \end{aligned}$$

Note that $p_B(x) = p(x)/p(B)$ for $x \in B$, and hence $r_B^i(x) = r^i(x) + \log(p(B)^{-1})$. Now Equation (7) implies

$$\forall x, x' \in B: |r_B^i(x) - r_B^i(x')| \leq 2R, \quad i = 1, \dots, t. \quad (8)$$

which finishes the proof. \square

It remains to observe that the covering number for our case is $\log N = \log N_0 + \log \log(|\mathcal{X}|) + \log \log(1/\epsilon) + \log(t/\delta)$ (see for example [Reg07]) which finishes the proof. \square

V. APPLICATIONS

Theorem 2 (Sampling preserves entropy rate [Vad03]). *Let \mathcal{X} be a fixed finite alphabet, and let X_1, \dots, X_t be a sequence of correlated random variables each over \mathcal{X} . Let $i_1, \dots, i_\ell \in [1, t]$, where $\ell < t$, be chosen from the set $\{1, \dots, t\}$ by an averaging (μ, θ, γ) -sampler. Then there is a random variable \mathcal{B} taking $m = t \cdot O(\log \log |\mathcal{X}| + \log \log(1/\epsilon) + \log t)$ bits, such that*

$$\begin{aligned} \frac{1}{\ell \log |\mathcal{X}|} \tilde{H}_\infty^\epsilon(X_{i_\ell} X_{i_{\ell-1}} X_{i_{\ell-2}} \dots, X_{i_1} | \mathcal{B}) &\geq \\ \frac{1}{t \log |\mathcal{X}|} \tilde{H}_\infty(X_{i_\ell} X_{i_{\ell-1}} X_{i_{\ell-2}} \dots, X_{i_1}) - \text{err}_{\text{Spoil}} - \text{err}_{\text{Sampler}} \end{aligned}$$

where the errors due to spoiling and sampling equal

$$\begin{aligned} \text{err}_{\text{Spoil}} &= O(\log \log |\mathcal{X}| + \log \log(1/\epsilon) + \log t) / \log |\mathcal{X}| \\ \text{err}_{\text{Samp}} &= O(\sqrt{\ell^{-1} \log(1/\epsilon)}). \end{aligned}$$

Remark 3 (Local extractors). *Composing this with an extractor over \mathcal{X}^ℓ one obtains a local extractor, which reads only a small fraction (specifically $\frac{\ell}{t}$) of input bits. We refer to [Vad03] for a general discussion.*

Proof. We will argue that the sequence $X_{i_1}, \dots, X_{i_\ell}$, for ℓ sufficiently big, likely has the same entropy rate (entropy per block) as the original sequence X_1, \dots, X_ℓ .

Let \mathcal{B} be the family guaranteed by Theorem 1. By part (a) of Corollary 1 applied t times (starting from $i = t$ downto $i = 1$) we have for every $B \in \mathcal{B}$

$$\sum_{i=1}^t H_\infty(X_i | X_{<i}, B) \geq H_\infty(X | B) - O(t\delta)$$

and now by part (b) applied to each summand

$$\sum_{i=1}^t \tilde{H}_\infty(X_i | X_{<i}, B) \geq H_\infty(X | B) - O(t\delta). \quad (9)$$

Note that

$$\mathbb{E}_{i_1, \dots, i_\ell} \sum_{j=1}^\ell \tilde{H}_\infty(X_{i_j} | X_{<i_j}, B) = \frac{1}{t} \sum_{i=1}^t \tilde{H}_\infty(X_i | X_{<i}, B)$$

In particular, with high probability over $(i_j)_{j=1, \dots, \ell}$

$$\frac{1}{\ell} \sum_{j=1}^\ell \tilde{H}_\infty(X_{i_j} | X_{<i_j}, B) \gtrsim \frac{1}{t} \sum_{i=1}^t \tilde{H}_\infty(X_i | X_{<i}, B). \quad (10)$$

For the sake of clarity, we comment later on the exact error in accuracy and probability in Equation (10). Observe that by part (b) of Corollary 1 we obtain

$$\begin{aligned} \frac{1}{\ell} \sum_{j=1}^\ell \tilde{H}_\infty(X_{i_j} | X_{i_j} X_{i_{j-1}} \dots, X_{i_1} | B) &\geq \\ \frac{1}{\ell} \sum_{j=1}^\ell \tilde{H}_\infty(X_{i_j} | X_{<i_j} | B) - O(\delta) \end{aligned} \quad (11)$$

(which is the *conditioning reduces entropy* property). Again, by applying part (b) of Corollary 1 to the sum on the right-hand side of Equation (11) we get

$$\frac{1}{\ell} \sum_{j=1}^{\ell} H_{\infty}(X_{i_j} | X_{i_j} X_{i_{j-1}} \dots, X_{i_1} | B) \geq \frac{1}{\ell} \sum_{j=1}^{\ell} \tilde{H}_{\infty}(X_{i_j} | X_{i_j} X_{i_{j-1}} \dots, X_{i_1} | B) - O(\delta) \quad (12)$$

By Lemma 3 from Equation (11) we get for every $B' \in \mathcal{B}'$

$$\frac{1}{\ell} H_{\infty}(X_{i_{\ell}} X_{i_{\ell-1}} X_{i_{\ell-2}} \dots, X_{i_1} | B) \geq \frac{1}{\ell} \sum_{j=1}^{\ell} \tilde{H}_{\infty}(X_{i_j} | X_{i_j} X_{i_{j-1}} \dots, X_{i_1} | B) - O(\delta). \quad (13)$$

Combining this with Equation (11), Equation (10) and Equation (9) we finally obtain (with high probability)

$$\frac{1}{\ell} H_{\infty}(X_{i_{\ell}} X_{i_{\ell-1}} X_{i_{\ell-2}} \dots, X_{i_1} | B) \gtrsim \frac{1}{t} H_{\infty}(X | B) - O(\delta). \quad (14)$$

Note that this holds for every B . Recall that by Remark 2 we can assume $\Pr[B] = \Omega(|\mathcal{B}|^{-1}\epsilon)$. Now by Lemma 4 we have $H_{\infty}(X | B) > H_{\infty}(X) - O(\log |\mathcal{B}| + \log(1/\epsilon))$ and thus

$$\frac{1}{\ell} H_{\infty}(X_{i_{\ell}} X_{i_{\ell-1}} X_{i_{\ell-2}} \dots, X_{i_1} | B) \gtrsim \frac{1}{t} H_{\infty}(X | B) - \frac{O(\log |\mathcal{B}| + \log(1/\epsilon))}{t}.$$

We can do slightly better. Namely, from Equation (14)

$$\frac{1}{\ell} \tilde{H}_{\infty}(X_{i_{\ell}} X_{i_{\ell-1}} X_{i_{\ell-2}} \dots, X_{i_1} | \mathcal{B}) \gtrsim \frac{1}{t} H_{\infty}(X | \mathcal{B}) - O(\delta). \quad (15)$$

where \mathcal{B} is a random variable that assigns to every point x the corresponding set B covering x , conditioned in addition on the map being defined (it fails when x is not covered by any B which happens w.p. at most ϵ) Indeed, we have

$$2^{-H_{\infty}(X_{i_{\ell}} X_{i_{\ell-1}} X_{i_{\ell-2}} \dots, X_{i_1} | B)} \leq \left(2^{-H_{\infty}(X_{i_{\ell}} X_{i_{\ell-1}} X_{i_{\ell-2}} \dots, X_{i_1} | \mathcal{B})} \right)^{\frac{\ell}{t}}.$$

and Equation (16) follows by averaging over B and the Jensen Inequality (note that $\ell < t$ implies that the corresponding mapping is concave). By Lemma 3 applied to $X | \mathcal{B} = B$ for all possible B to outcomes of \mathcal{B} we obtain

$$\frac{1}{\ell} \tilde{H}_{\infty}(X_{i_{\ell}} X_{i_{\ell-1}} X_{i_{\ell-2}} \dots, X_{i_1} | \mathcal{B}) \gtrsim \frac{1}{t} H_{\infty}(X) - \frac{\log |\mathcal{B}|}{t} - O(\delta) \quad (16)$$

It remains to use an explicit bound on $|\mathcal{B}|$ from Theorem 1, set the sampler to $\gamma = \epsilon$ and compute θ from ϵ and ℓ . \square

VI. CONCLUSION

By a simple combinatorial argument combined with the spoiling knowledge technique we showed how to exhibit strong block-source structures in any min-entropy source. This approach may be applied to locally enforce chain rules (or other desired properties) for min-entropy.

REFERENCES

- [BBCM95] C. H. Bennett, G. Brassard, C. Crépeau, and U. M. Maurer. “Generalized privacy amplification”. In: *IEEE Trans. Information Theory* (1995).
- [BKR16] M. Bellare, D. Kane, and P. Rogaway. “Big-Key Symmetric Encryption: Resisting Key Exfiltration”. In: *CRYPTO*. 2016.
- [BRSW06] B. Barak, A. Rao, R. Shaltiel, and A. Wigderson. “2source Dispersers for Sub-polynomial Entropy and Ramsey Graphs Beating the Frankl-Wilson Construction”. In: *STOC '06*. 2006.
- [BST03] B. Barak, R. Shaltiel, and E. Tromer. “True Random Number Generators Secure in a Changing Environment”. In: *CHES*. 2003.
- [Cac97] C. Cachin. “Smooth Entropy and Rényi Entropy”. In: *EUROCRYPT '97*. 1997.
- [DFRSS07] I. Damgård, S. Fehr, R. Renner, L. Salvail, and C. Schaffner. “A Tight High-Order Entropic Quantum Uncertainty Relation with Applications”. In: *CRYPTO*. 2007.
- [DORS08] Y. Dodis, R. Ostrovsky, L. Reyzin, and A. D. Smith. “Fuzzy Extractors: How to Generate Strong Keys from Biometrics and Other Noisy Data”. In: *SIAM J. Comput.* (2008).
- [DP08] S. Dziembowski and K. Pietrzak. “Leakage-Resilient Cryptography”. In: *FOCS*. 2008.
- [IS13] M. Iwamoto and J. Shikata. “Information Theoretic Security for Encryption Based on Conditional Rényi Entropies”. In: *ICITS*. 2013.
- [NZ96] N. Nisan and D. Zuckerman. “Randomness is Linear in Space”. In: *J. Comput. Syst. Sci.* (1996).
- [Reg07] L. C. Rego. *Probability Simplex Covering number*. <http://www.de.ufpe.br/~leandro/Coveringnumber.pdf>. 2007.
- [RW04] R. Renner and S. Wolf. “Smooth Renyi entropy and applications”. In: *ISIT*. 2004.
- [Sha11] R. Shaltiel. “An Introduction to Randomness Extractors”. In: *ICALP*. 2011.
- [Vad03] S. P. Vadhan. “On Constructing Locally Computable Extractors and Cryptosystems in the Bounded Storage Model”. In: *CRYPTO*. 2003.